

# Extending Kubernetes or: How I learned to Stop Worrying and Trust the Spec

---

David Zhu

# About Me

- Software Engineer @ Google
- SIG-Storage Contributor
- Working on/with Container Storage Interface (CSI) Specification since pre-v0.1.0
  - Released v1.0.0 last week
- Developer for the GCP Persistent Disk CSI Driver

# About Me

- Recently broke the record for the Houdini Escape Room in SF



# Kubernetes Storage circa v1.9



Being a  
maintainer was  
hard

**“Code for storage plugins that are not supported on your platform is still a risk”**

- Michelle Au, Kubernetes Storage Guru

Being a  
storage vendor  
was hard

**“ If we found a critical bug in our plugin after the Kubernetes release, we don’t have a way to release the plugin separately.”**

- Xing Yang, OpenSDS Developer

# Some of the Problems to solve

1. Allow vendors to iterate at their own pace
  - a. Release Decoupling
2. Stop crashes in drivers bringing down all of Kubernetes
  - a. Binary Decoupling
3. Reducing maintenance burden for Kubernetes Approvers
  - a. Code Decoupling

# The Container Storage Interface



# What is Container Storage Interface (CSI)?

- A few sidecar containers
- Drivers
- Some pieces inside the container orchestrator (such as Kubernetes)
- Most importantly... The Specification



CONTAINER  
STORAGE  
INTERFACE



A Specification  
Needs an  
Ecosystem and  
Process

# Changing the Spec

—

# Concurrent Development and Understanding



MESOS



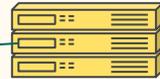
CLOUDFOUNDRY



CONTAINER  
STORAGE  
INTERFACE



portworx



# Transparent & Open Development Forum



# NodeStageVolume/NodeUnstageVolume #169

Edit

**Merged** saad-ali merged 2 commits into container-storage-interface:master from davidz627:feature/MountDevice on Feb 20

Conversation 114   Commits 2   Checks 0   Files changed 3

+732 -174



davidz627 commented on Jan 17 • edited

Contributor + 👤 ...

Fixes Issue: #119

Adds two new RPCs, NodeStageVolume/NodeUnstageVolume. These RPCs MUST be called by the CO if the Node Plugin advertises the STAGE\_UNSTAGE\_VOLUME capability. Plugins that advertise this capability SHOULD defer volume reference counting to the CO.

This PR supersedes #122, @akutz gets credit for the majority of the code, I'm just picking up where he left off and pushing this through.

cc @saad-ali @jieyu @clintkitson @def

Reviewers

- akutz ✓
- cpuguy83 ✓
- jieyu ✓
- jdef ✓
- saad-ali ✓

Assignees

# Add Snapshot Support in CSI Spec #207

**Closed** xing-yang wants to merge 1 commit into container-storage-interface:master from xing-yang:snapshot

Conversation 248 Commits 1 Checks 0 Files changed 3

+1,410 -471



xing-yang commented on Mar 1 • edited ▾

Contributor + 👤 ...

This PR proposes to add snapshot support in CSI specification. It proposes to add 3 new controller service RPCs, including CreateSnapshot, DeleteSnapshot, and ListSnapshots, and modify the existing ControllerGetCapabilitiesResponse and CreateVolumeRequest.

There are several cases for using the snapshot data in a volume:

1. Create volume from snapshot: In this case, a new volume will be created on the storage backend and the data in the snapshot will be copied to the new volume. This case will be addressed by this proposal.
2. Revert snapshot: In this case, data in the original volume will be erased and replaced with data in the snapshot. This case is out of scope for this proposal and should be handled by a function of its own (not to be mixed with create volume from snapshot) if it is proposed in the future.

The following changes will be made in the Service Controller:

- Add CreateSnapshot

Reviewers

- |  |              |   |
|--|--------------|---|
|  | iliastsi     | 🗨 |
|  | jsafrane     | 🗨 |
|  | jdef         | 🗨 |
|  | yasker       | 🗨 |
|  | humblec      | 🗨 |
|  | Akrog        | ✕ |
|  | gouthampacha | 🗨 |
|  | saad-ali     | 🗨 |
|  | lichuqiang   | 🗨 |
|  | e4jet        | 🗨 |
|  | govint       | 🗨 |
|  | jieyu        | ✕ |
|  | julian-hj    | ✕ |
|  | cpuguy83     | ✓ |
|  | j-griffith   | ✓ |



**xing-yang** commented on May 16

Contributor



Just want to drop a note here that PR [#224](#) (which I submitted to replace this original PR because some people can no longer open and review this PR) was merged today. Thank you all for reviewing and providing valuable feedback!

Open

Collaboration

Promotes

Consensus

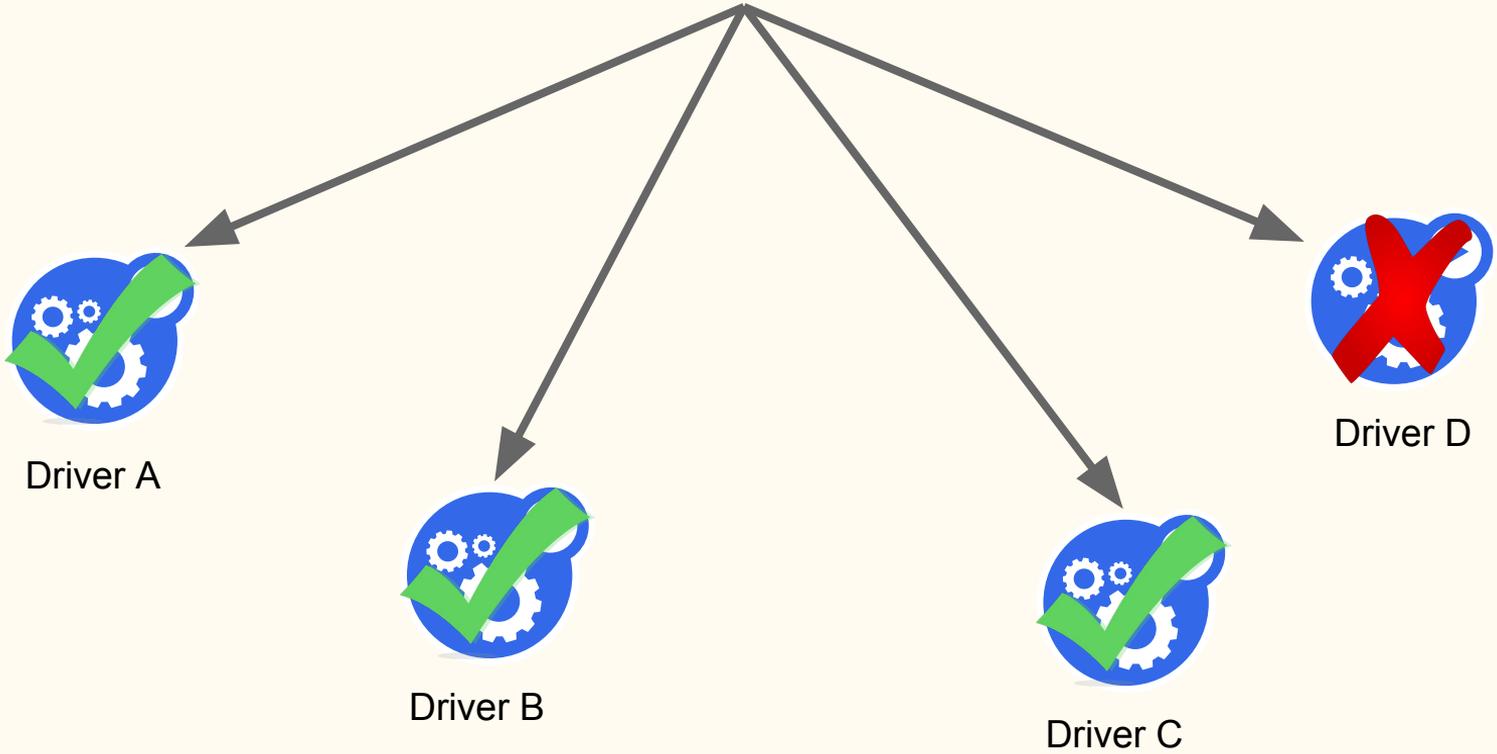
Using the Spec

—

# Different Interpretations



# CSI Sanity



```
• Failure [0.001 seconds]
Controller Service
/usr/local/google/home/dyzz/go/src/sigs.k8s.io/gcp-compute-persistent-disk-csi-driver/vendor/github.com/kubernetes-csi/csi-test/pkg/sanity/te
sts.go:44
  ControllerPublishVolume
  /usr/local/google/home/dyzz/go/src/sigs.k8s.io/gcp-compute-persistent-disk-csi-driver/vendor/github.com/kubernetes-csi/csi-test/pkg/sanity
/controller.go:712
    should fail when the volume does not exist [It]
    /usr/local/google/home/dyzz/go/src/sigs.k8s.io/gcp-compute-persistent-disk-csi-driver/vendor/github.com/kubernetes-csi/csi-test/pkg/sani
ty/controller.go:856

    Expected
    <codes.Code>: 3
    to equal
    <codes.Code>: 5

    /usr/local/google/home/dyzz/go/src/sigs.k8s.io/gcp-compute-persistent-disk-csi-driver/vendor/github.com/kubernetes-csi/csi-test/pkg/sani
ty/controller.go:882
-----
E1206 13:00:21.780099 144829 utils.go:55] GRPC error: rpc error: code = NotFound desc = could not split nodeID: failed to get id components
. expected projects/{project}/zones/{zone}/instances/{name}. Got: some-fake-node-id
E1206 13:00:21.782340 144829 utils.go:55] GRPC error: rpc error: code = AlreadyExists desc = Disk sanity-controller-published-incompatible
-FA819416-883F3F2E already published to node projects/test-project/zones/country-region-zone/instances/test-name but incompatible: disk mode
does not match. Got READ_WRITE. Want READ_ONLY
E1206 13:00:21.783225 144829 utils.go:55] GRPC error: rpc error: code = InvalidArgument desc = ControllerUnpublishVolume Volume ID must be
provided
..

Summarizing 1 Failure:

[Fail] Controller Service ControllerPublishVolume [It] should fail when the volume does not exist
/usr/local/google/home/dyzz/go/src/sigs.k8s.io/gcp-compute-persistent-disk-csi-driver/vendor/github.com/kubernetes-csi/csi-test/pkg/sanity/c
ontroller.go:882

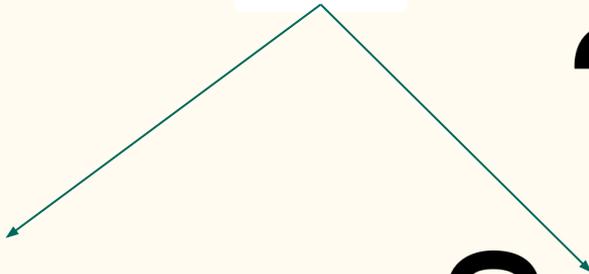
Ran 55 of 57 Specs in 0.042 seconds
FAIL! -- 54 Passed | 1 Failed | 0 Pending | 2 Skipped --- FAIL: TestSanity (0.05s)
```

**NotFound !=  
InvalidArgument**

Conformance  
Tools  
Maintain  
Consistency

# Getting Adoption





In-Tree Plugin



CONTAINER  
STORAGE  
INTERFACE



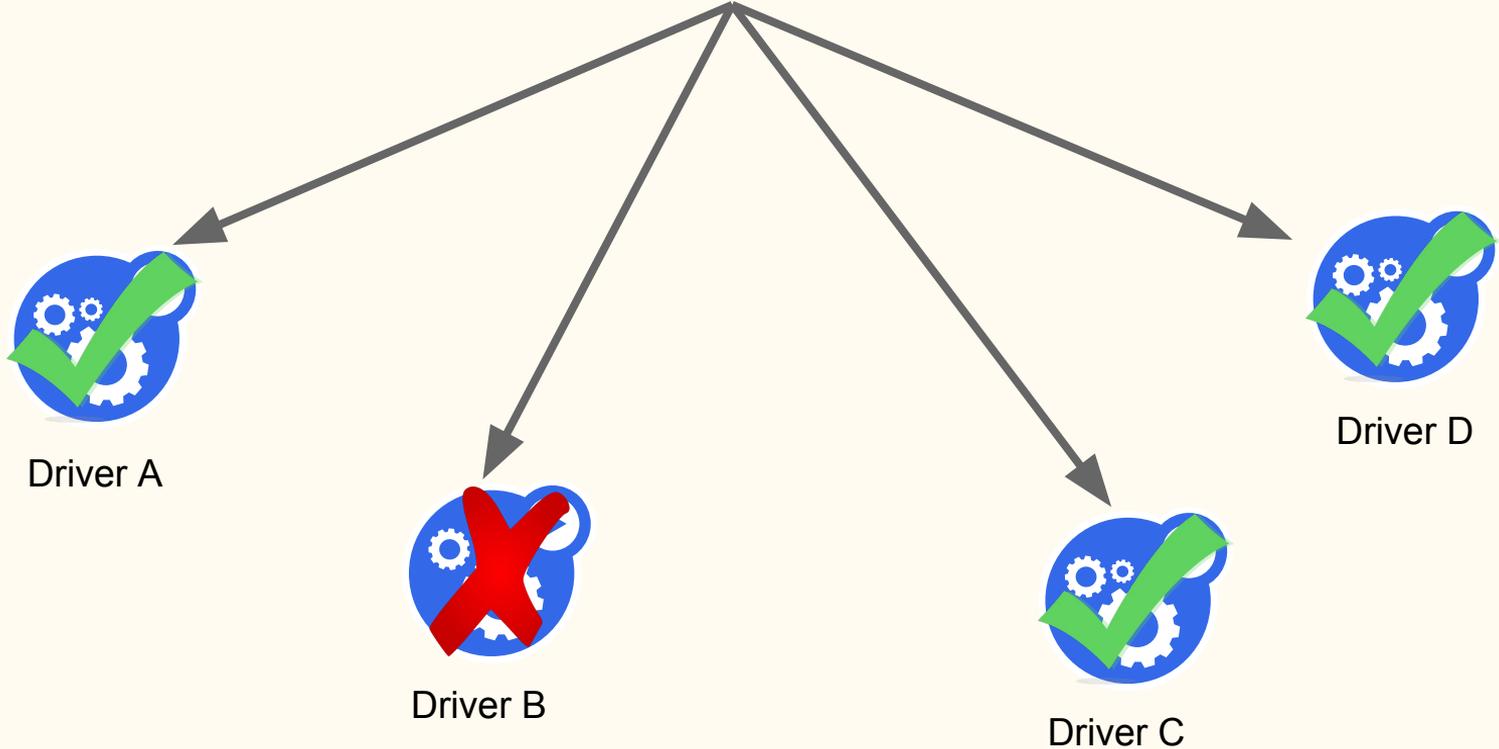
# User Traction Through Value Add



# Volume Snapshots



# Future: Testing Framework



Adding  
Value  
Accelerates  
Adoption

# Extending to Extra Extensions

—

Open

Collaboration

Promotes

Consensus

Conformance  
Tools  
Maintain  
Consistency

Adding  
Value  
Accelerates  
Adoption

The  
Specification  
Isn't  
Everything

Extending Kubernetes or: How I  
learned to

~~Stop Worrying and Trust the Spec~~

Build a Consistent Ecosystem and

Process Around a Specification

---

David Zhu, Software Engineer @ Google

Email: [dyzz@google.com](mailto:dyzz@google.com)

Github: davidz627